



11 Publication number : **0 567 294 A2**

12

EUROPEAN PATENT APPLICATION

21 Application number : **93303049.6**

51 Int. Cl.⁵ : **H04M 3/50, H04M 7/00**

22 Date of filing : **20.04.93**

30 Priority : **21.04.92 US 871770**

43 Date of publication of application :
27.10.93 Bulletin 93/43

84 Designated Contracting States :
BE DE DK ES FR GB GR IT LU NL SE

71 Applicant : **BOSTON TECHNOLOGY INC.**
100 Quannapowitt Parkway
Wakefield, Massachusetts 01880 (US)

72 Inventor : **Arumainayagam, Allen Theivendran**
30 Daniels Street No. 504
Malden, MA 02148 (US)
Inventor : **Penfield, Robert Flagg**
22 Howard Road
Maynard, MA 01754 (US)
Inventor : **Capo, José**
20 Windham Road
Pelham, NH 03076-2370 (US)
Inventor : **Deigh, Anthony Nicholas**
601 Townsend Road
Groton, MA 01450 (US)

74 Representative : **Skone James, Robert**
Edmund
GILL JENNINGS & EVERY, Broadgate House,
7 Eldon Street
London EC2M 7LH (GB)

54 **Multi-system network addressing.**

57 Storage locations, such as voice mailboxes, are addressed on a plurality of widely distributed host systems, connect via a network, by defining logical domains each of which may include storage locations on more than one host system. The domains are defined by address tables which indicate which host systems provide ranges of address locations. Transfer of data between domains is simplified by the use of translation tables. one type of translation table is an alias table which has entries containing a range in the originating domain, a new domain and a starting point for a range in the new domain, thus defining a one-to-one correspondence between the ranges in the two domains. Another translation table is a gateway table which defines a gateway to a new domain when an input address has characteristics matching a prefix, number of digits, or both, of an entry in the gateway table. Upon finding the closest possible match in the gateway table, the portion of the input address after the prefix, if any, is used as the address in the new domain.

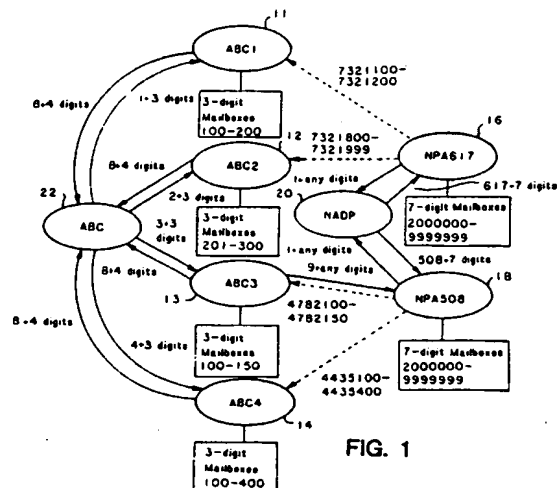


FIG. 1

EP 0 567 294 A2

ings:-

Figure 1 is a block diagram of domains used for addressing according to the present invention;

Figure 2 is a block diagram of a data processing system for providing information services which may use addressing according to the present invention;

Figure 3 is a flow chart of a method according to the present invention; and,

Figure 4 is a module visibility diagram of software that may be used to implement the present invention.

In existing telephone networks, dialling a series of digits on one communication device, such as a telephone or facsimile machine, identifies another communication device at a specific physical location. The digits dialled inform the telephone network of the physical location of the called device and switches in the network connect the calling device to that physical location. The network used may be the public telephone network, largely provided by local Bell operating companies in the United States, or a private network or a leased line direct from one location to another location of a single large corporation.

Newer information services which are being made available by telephone companies do not require the same physical relationships as conventional telephone devices. For example, a voicemail subscriber may have a voice mailbox provided by a data processing system hundreds or thousands of meters from the telephone whose calls are routed to that voice mailbox. Furthermore, the subscriber may be able to access that mailbox from any telephone, not only the telephone whose calls are routed to that voice mailbox.

As the number of subscribers to services, such as voicemail, increase, an efficient widely distributed network of data processing systems is highly desirable. To enable a plurality of widely distributed data processing systems to provide information services, transfer of data from one system to another generally requires the actual transfer of the data. Furthermore, efficient use of resources may mean that two telephones sitting on adjacent desks may have voice mailboxes associated therewith provided by data processing systems which may be several kilometres apart and several kilometres from the telephones. Presumably, the adjacent telephones have a logical relationship, that is, the owner of the telephones is a company which may have 100 or more telephones nearby the two adjacent telephones and hundreds or thousands more at other sites. The voice mailboxes associated with all of these telephones thus have a logical relationship, i.e., they are subscribed for by the same company. According to the present invention, this logical relationship is termed a domain or logical domain.

As an example, Figure 1 depicts the domains of mailboxes of the ABC company with less than a thousand voice mailboxes in four logical domains 11-14. Other logical domains in Figure 1 are represented by ovals 16, 18 for voice mailboxes assigned to telephone numbers in numbering plan area (NPA) codes 617 and 508, respectively. The NPA domains 16 and 18 serve as parent domains for the ABC domains 11-14. An oval 20 representing a North American dialling plan (NADP) routing domain provides a central place to define relationships to all area code domains, but does not include any voice mail boxes. Similarly, oval 22 represents a routing domain for the entire ABC company which enables voice mail messages to be sent from, e.g., the ABC1 domain 11 to the ABC4 domain 14 by outdialling "84nnn" where "nnn" represent any number in the range 100-400. In a similar manner, a subscriber in the NPA508 domain 18 uses the NADP routing domain 20. Further explanation of how voice mail messages are handled using logical domains will be provided below.

Voice mailboxes and other types of storage locations of course have a physical manifestation in addition to the logical relationships illustrated in Figure 1. In the case of voice mailboxes, the voice data or messages must actually be stored in some physical device. The type of device will depend upon the type of data processing system being used, and will be referred to as simply a storage unit. As known in the art, the storage unit may be provided by magnetic disks, optical data storage, etc. The storage units are connected to data processing systems which will be referred to as host systems, or simply "hosts". A relationship between the logical storage locations and the physical devices providing storage for those locations is defined according to the present invention in an address table. One example of an address table is provided below for a portion of the NPA617 domain 16. A range of address locations having a scope, or number of locations, of at least one, is defined by the first two columns of the address table. The third column contains a host identifier of the host system providing physical storage for the storage locations or voice mailboxes within the range. Alternative ways of defining the address table include using a starting point of a range and length of the range and any other known method of identifying storage locations provided by a host system.

from either NPA domain 16, 18 the NADP routing domain 20 is accessed. To permit access to other domains using 10 or more digits from within the NPA domains 16, 18, the closest match is always chosen first. Thus, if there is an entry in a gateway table with a prefix of more than 7 digits, that entry will be checked first, followed by any entries with 7 digit prefixes, 6 digit prefixes, etc. Only then is access available using the entries in the gateway table which do not have prefixes, such as an entry for the NADP routing domain from the NPA617 domain 16. When no matches are found in the gateway table, an alias table is checked, as described below.

The second type of translation table is an alias table. The alias table maps addresses in one domain to addresses in another domain with a one-to-one correspondence. As indicated in the example, a range of several address locations can be mapped, but the range may also have a scope of one, e.g., 7321201-7321201. The address codes defining the beginning and end of the range are used to translate an input address into an addressed location in another domain. A portion of the alias table for the NPA617 domain 16 is provided below.

ALIAS TABLE FOR NPA617			
Start of Range	End of Range	Domain	Start Address
7321100	7321200	ABC1	100
7321800	7321999	ABC2	201

The primary use of aliases is to define groups of voice mailboxes associated with CENTREX lines. Within an area code different companies may have CENTREX phone numbers with the same first three or four digits. Within a CENTREX group, e.g., ABC1, only the last three (or four) digits are dialed to access a telephone or its associated voice mailbox, but to reach the same voice mailbox from the parent domain, e.g., NPA617, all seven digits are required. If a voice mailbox or other application account is outside an originating domain having three digit mailboxes, a three digit number may be used to identify another domain in an alias table. For example, the ABC1 domain 11 may have an alias table identifying a range between 400 and 499 as a range of addresses of voice mailboxes in XYZ Company.

A host system according to the present invention may be constructed as illustrated in Figure 2. Control of the system is provided by a processor (CPU) 58 which is connected via a passive blackplane 63 to hard drives 64, 66 via a disk controller 68. The host system is connected via digital switches 74 and a T-1 interface 76 to a telephone network. The storage locations or voice mailboxes are physically provided by the hard drive units 64, 66 and the translation tables and address table are also stored on these hard drive units 64, 66. The CPU 58 is programmed to provide services to subscribers including identifying an address location for each input address based on the tables stored in the hard drives 64, 66. The system illustrated in Figure 2 may be only a single processor system or a locally distributed host system having a plurality of such processor systems, as disclosed in more detail in U.S. Patent 5,029,199 and U.S. Patent Application Serial No. 07/594,648.

An overview of how the processor 58 in the system illustrated in Figure 2 is programmed to operate is illustrated in Figure 3. As described above, the physical manifestation of the domains and the relationships between the domains are defined 80 by storing address, alias and gateway tables. When an input address is received 82, the entries in the gateway table are checked 84, starting with the longest prefixes and ending with the shortest entry containing a number of digits only. If the input address matches 84 the characteristics of one of the entries in the gateway table, the input address is translated 86 into a new domain via the gateway defined in the table. If there is no match 84 of a gateway characteristic, the input address is compared 88 with the entries in the alias table. If the input address fits in one of the ranges defined in the alias table for the originating domain, the input address is translated 90 into an addressed location in an alias domain. In either case, translation then continues in the new domain. The storage location identified by the remainder of the input address, which would be all of the input address if not translated via the alias or gateway tables, is accessed 92 to perform the service requested by the user.

A more detailed description of the software used in the preferred embodiment will be provided with reference to Figure 4. A module visibility diagram is illustrated in Figure 4 for the software executed by the system illustrated in Figure 2. Each bubble in the diagram illustrated in Figure 4 represents a software object, i.e., procedure(s) and data related to each other. The higher level objects are at the top of Figure 4 and the lower level objects are at the bottom of Figure 4.

The highest level objects illustrated in Figure 4 are addressing domain administration program 102 and applications 104. The addressing domain administration program 102 is used to create and maintain the address, alias and gateway tables discussed previously. The corresponding name for the addressing domain administration program is adadmin. The applications object 104 represents all of the applications, i.e., programs, such as voicemail, which use these three tables to locate a mailbox address on a host system. The applications

2. A method of addressing in a distributed data processing system having storage units and processors, comprising the steps of:
 - (a) defining domains of storage locations in the storage units without requiring any physical relationship between the domains and the storage units containing the storage locations;
 - (b) storing at least one translation table to identify the domains for input addresses; and
 - (c) identifying an address location for each of the input addresses by determining an addressed domain and then the addressed location within the addressed domain.
3. A method according to claim 1 or 2, wherein the storage locations include application accounts and the distributed data processing system provides information services to subscribers using at least two computer systems separated by at least 100 meters and connected by a telephone switching system.
4. A method according to claim 1, 2 or 3, wherein the storage is formed of host systems each including at least one of the storage units and at least one of the processors, and
 - wherein said defining in step (a) comprises the step of (a1) storing an address table for each logical domain included in the domains in at least one of the storage units accessible to all of the processors, the address table defining physical locations of the storage locations within the logical domain, such that each of the storage locations is defined as located on one of the host systems and more than one of the host systems may provide the storage locations within the logical domain.
5. A method according to claim 4,
 - wherein said storing in step (a1) includes storing a range of the address locations having a scope of at least one and a host identifier for each entry in the address table, and
 - wherein said defining in step (a) further comprises the step of (a2) copying the address table for each of the domains to all of the host systems.
6. A method according to claim 4 or 5,
 - further comprising the step of (b) receiving the input addresses to identify the storage locations to be accessed,
 - wherein said storing in step (d) comprises the step of storing in at least one of the storage units entries in an alias table providing one of the at least one translation table for the logical domain included in the domains, each entry in the alias table defining a one-to-one correspondence between a range of address codes and a range of address locations in one of the domains, the ranges of address codes and address locations each having a scope of at least one, and
 - wherein said identifying in step (c) comprises the steps of:
 - (c1) comparing the input addresses with the address codes in the alias table for the logical domain; and
 - (c2) identifying the addressed locations as one of the storage locations having a one-to-one correspondence with a selected address code in the range of address codes in one of the entries of the alias table if the selected address code matches one of the input addresses.
7. A method according to claim 4 or 5,
 - further comprising the step of (d) receiving the input addresses to identify the storage locations to be accessed, wherein said storing in step (b) comprises the step of storing, in at least one of the storage units, entries in a gateway table providing one of the at least one translation table for the logical domain included in the domains, each entry defining a relationship between characteristics of the input addresses and one of the domains, and
 - wherein said identifying in step (c) comprises the steps of:
 - (c1) comparing the input addresses with each entry in the gateway table; and
 - (c2) identifying the address location as within an addressed domain in one of the entries in the gateway table for each of the input addresses having a match with the characteristics in the one of the entries.
8. A method according to claim 7,
 - wherein said storing in step (b) comprises the step of (b1) storing prefixes of the input addresses as one of the characteristics, and
 - wherein said identifying in step (c) identifies the addressed location as within the addressed domain if one of the prefixes matches an initial part of one of the input addresses.
9. A method according to claim 8, wherein said storing in step (b) further comprises the step of (b2) storing

an address table defining domains of the storage locations in said storage units without requiring any physical relationship between the domains and said storage units containing the storage locations within the domains and at least one translation table to identify the domains for input addresses; and

5 processors (58) to provide services to subscribers of said distributed data processing system and to identify an address location for each of the input addresses by determining an addressed domain and then the addressed location within the addressed domain.

10

15

20

25

30

35


40

45

50

55

FIG. 2



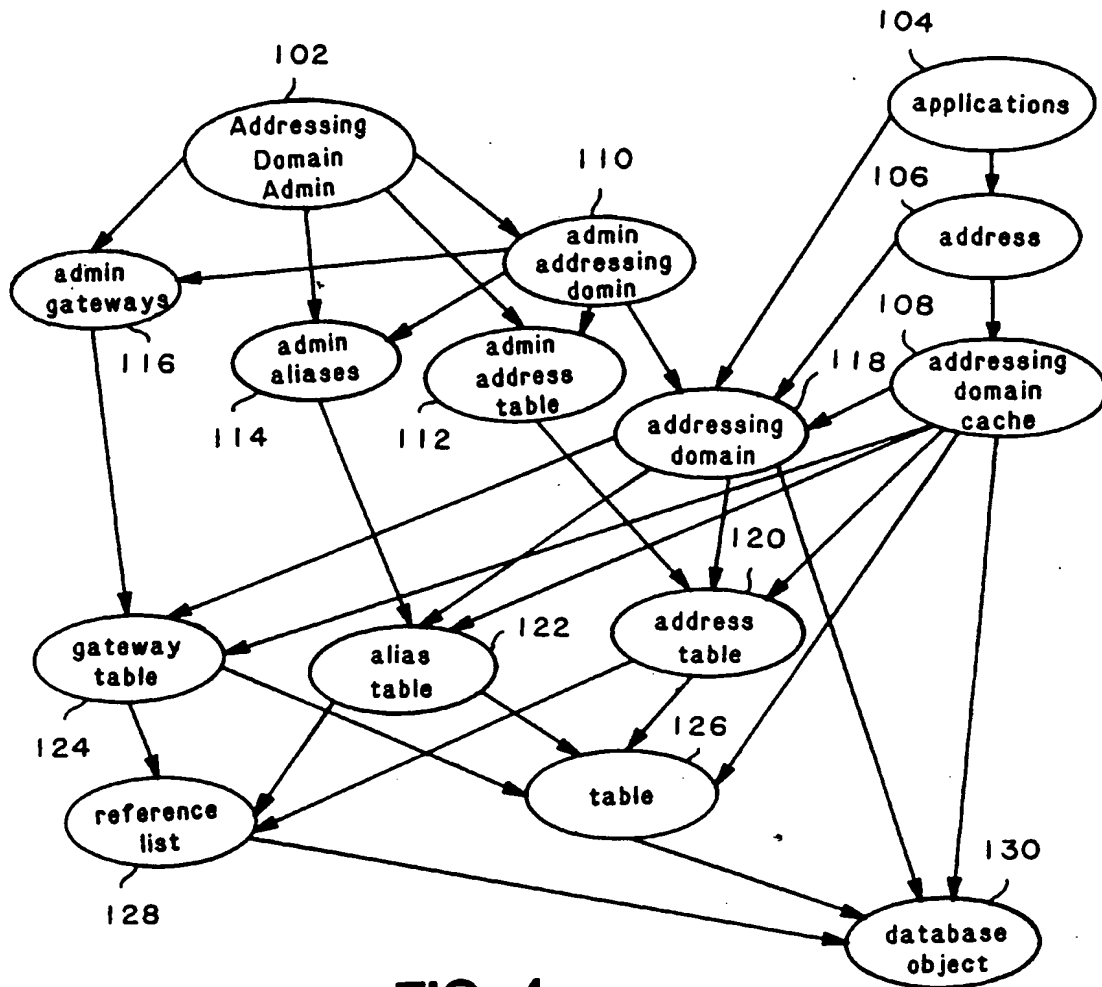


FIG. 4